

# PERMANENT TCP CONNECTIONS ACROSS SYSTEM REBOOTS

## BACKGROUND OF THE INVENTION

### Technical Field

5 This invention relates to maintaining connections for both mobile and non-mobile nodes in computer networks.

## DESCRIPTION OF THE PRIOR ART

10 Most mobile devices, such as laptop computers, are frequently shutdown for a variety of reasons. For instance, most are battery-powered, so they often shutdown to conserve battery power. Airlines require such devices to be turned off when taking off or landing regardless of the power source. However, each shutdown causes termination of all 15 applications and a loss of network connectivity, which results in substantial inconvenience to the user.

15 A problem resulting from the shutdown of a mobile device during an application can be illustrated by an example in which a traveler using a laptop computer in an airport connects to a database server and enters a complex query, which would normally take a long time to execute. However, when the traveler boards the plane, the system must be shutdown prior to takeoff. Upon reaching his destination, the traveler restarts the computer, rebooting the system and restarting the application. Yet, without permanent TCP connections across a 20 system boot, a new connection must be established and the query must be reissued before the reply can be received. Accordingly, prior to the present invention there has been a need for a method by which mobile IP can remain connected across reboots.

## SUMMARY OF THE INVENTION

25 The invention provides a method, system and apparatus that enables clients to keep their connections open during reboots or shutdowns of networked or mobile computer systems. The invention allows an end-user to perform orderly system shutdowns of mobile systems (especially useful to save battery power on mobile systems), without risk of losing 30 transactions on open TCP connections. The invention further suspends each ongoing transaction while the client system is down, and resumes it when the system comes back up without loss of connectivity.

The mobile device of the present invention relies upon 'mobile IP' (RFC 2002) for connectivity. Mobile IP provides a static, unchanging IP address to the remote communication endpoint. The constant address allows the higher layer, such as TCP, to work with the remote endpoint as it travels from one location to another without loss of data.

5 The present invention provides the additional advantage of retaining connectivity, even though the device is shutdown.

Other features and advantages of this invention will become apparent from the following detailed description of the presently preferred embodiment of the invention, taken in conjunction with the accompanying drawings.

## BRIEF DESCRIPTION OF THE DRAWINGS

FIG. 1 is a flowchart showing the preferred embodiment of the invention.

FIG. 2 is a flowchart showing placement of a remote endpoint in a persist state, as needed in the preferred embodiment of the present invention

15 FIG. 3 is a flowchart showing an alternative method for putting the remote endpoint in a persist state.

FIG. 4 is a flowchart showing a method for rebooting the system after normal shutdown.

## DESCRIPTION OF THE PREFERRED EMBODIMENT

### Overview

In general, FIG. 1 depicts the process for shutting down the system according to the preferred embodiment without loss of connectivity.

### 25 Technical Background

#### Configuration

A system utility is used in the invention to mark the TCP endpoints that are maintained across reboots, step 12. The endpoints are described by the local port, local address, foreign port and foreign address. Alternatively, a method is provided by which 30 endpoints that need to survive the shutdown/reboot cycle may be selected at the time of system shutdown.

## Orderly Shutdown of the TCP Endpoint (Deactivation)

“Orderly” shutdown, is a deliberate action on the part of the user, as opposed to a system crash or failure, or a power cycle. In the invention, when the system shutdown command is given to the operating system, step 13, it alerts the TCP engine. The system waits for the TCP engine to save the TCP state, for the endpoints configured to survive the shutdown and reboot cycle. If the shutdown is a panic shutdown, the shutdown is handled normally. If the shutdown is orderly, the following steps are performed simultaneously.

The TCP endpoint stops sending or accepting data from the application, step 14. The TCP engine fails or blocks all read/write calls from the application. An implementation might fail the command with a suitable error code to let the application know that the connection is being ‘deactivated.’ Since the device is ‘going down,’ all interfaces are blocked from accepting new data.

The local endpoint acknowledges all data that is in the ‘receive’ buffer, and advertises a window of 0 (zero) in the acknowledgement. The remote endpoint then enters a TCP persist state.

In the preferred embodiment, the state of the TCP endpoint is stored on the disk or software agent, step 16. The state includes such information as the source and destination address, the port numbers, the window sizes, maximum size, data including any urgent data, and other information needed to maintain a TCP connection.

Received data that has been acknowledged, but not passed to the application, is stored with the TCP state. The data that was received from the application, but not sent across, or sent across but not acknowledged by the other end, is also stored with the TCP state.

A unique ID, along with the application name, is also stored with the TCP state. An implementation may store other information needed to identify the endpoint in accordance with the needs of the operating system and the environment in which the application is running.

Once the TCP state has been saved, the socket and TCP control block are dissociated from the file descriptor, step 17, to avoid the endpoint from being closed when the application exits and the file descriptor is closed. In the alternative, if the socket and TCP

control block are not disassociated, the TCP engine sends a termination indication (FIN) to the other endpoint.

### **Orderly Shutdown of the Application**

5 At the time the TCP state is saved, the application state is also saved, step 18, although the application on a mobile device is typically lightweight, *i.e.*, without much state.

Alternatively, the application may be allowed to run to some defined ‘sync’ point, and then the state is saved. There is no need to repeatedly check the state.

### **10 System Reboot and Application Restart**

Figure 4 shows the process for rebooting the system according to the preferred embodiment. At the time the system reboots, a utility, referred to as the “TCP-reacquire daemon,” is started to read in all of the saved TCP endpoint states, step 43. These endpoint states are added to the port/IP address table of the TCP engine using a system call. The endpoints are marked, so that data on the endpoints will not be accepted/sent until an application actually acquires the endpoint. It also sets up the data structures that are required to manage the connection(s) and initializes the structures with the saved data.

15 The application is then restarted, step 44. As noted above, the application is expected to either recover from the saved ‘sync’ point, or be stateless. The application asks for all of 20 the endpoints that were previously connected to it, by application name or ID, step 45.

Alternatively, the application may request a connection to the remote endpoint by specifying the remote address and port.

### **Reactivating the Connection**

25 Reactivation of the connection, step 46, is supported by modifications internal to the operating system kernel, as described below. The TCP engine endpoint is accessed by the application using the existing network APIs (Application Program Interface). For example, in the preferred embodiment, the BSD sockets API is used, although other network APIs may be modified in the same manner, as necessary.

30 In practice, to reactivate the connection, the application initially creates the communication endpoint, *e.g.*, a socket() system call. For example, the socket API bind()

call would discover if the requested binding already exists. If so, the binding would have been marked to indicate that it is a 'reacquire endpoint' by the 'TCP -reacquire daemon.' In the case in which connect() is called with an implicit bind(), the lookup table is searched with the foreign address and port to discover any existing bindings that can be reacquired.

5 The socket is linked with the TCP binding and the data associated with the binding is linked with the socket when the connect() call is made. The TCP control block is then 'unmarked' to allow it to send/receive data. The remote endpoint is sent a byte window advertisement that is greater than zero (window is opened). This causes the remote peer to come out of persist state, and the two endpoints can then exchange data normally.

10

### **Causing the Remote Endpoint to Keep the Connection**

In the absence of the solution presented in the invention, the remote endpoint can be in one of the following three states when the mobile node shuts down ("node" is a generic term referring to individual hardware components that make up a network, *e.g.*, general-purpose computers, stand-alone terminals, portable computers, servers, switches, routers and the like).

15 i) The endpoint has transmitted all data it had and all data has been acknowledged. The remote endpoint will not transmit any data since it has none to send. It might send KEEPALIVE probes, if it has been so configured. If the remote endpoint receives some data 20 from its application that needs to be sent to the mobile node, the remote endpoint will attempt transmission, which leads to situation (ii) described below.

25 ii) The endpoint has sent some data that has not been acknowledged. The remote endpoint will attempt to retransmit if it does not receive acknowledgement for this data. However, after certain number of attempts without acknowledgment, the remote endpoint will terminate the connection.

iii) The remote endpoint is in persist state, wherein the mobile node has closed the TCP stream window, in which case the remote host will periodically poll the mobile endpoint. If no response is received, the connection will be terminated.

30 Accordingly, all three situations cause the connection to terminate if the remote server does not receive a response. The remote endpoint sets an upper limit on the number of

tries that will be made, after which, upon not receiving an answer, the connection will be terminated.

Alternatively, as shown in FIGs. 2 and 3, the remote endpoint is made to keep the connection. In both cases, the alternatives cause the remote endpoint to enter a persist state.

5

### **Enhancement to Mobile IP (RFC 2002)**

A home “agent” is contacted by the mobile node (whether in the home network or in foreign network), and informed of its intention of shutting down. The mobile node requests the home agent to handle this connection. It sends a zero byte window advertisement to the remote endpoint to place it into persist mode, step 23. The node then requests the agent to handle the persist probes from the remote endpoint. Upon receiving the window advertisement, the remote node will enter persist state and send persist probes, periodically and repetitively polling the connection, step 26.

Whenever the home agent receives a packet from the remote host, as in step 26, it responds with the window advertisement of 0 (zero bytes), step 27. The window advertisement is based on the information it had recorded when the mobile node had sent the request packet.

If the remote endpoint did not receive the packet, it might send a data packet. In response to any message from the remote endpoint, the agent is requested to respond with a window advertisement of 0, on behalf of the mobile node. The agent is given the exact packet that it needs to send along with the address/port of the remote endpoint and the address and port of the mobile node. The agent records this information and responds to the mobile node with an acknowledgment, step 28.

Alternately, the agent could be given the details of the response packet, instead of the exact response packet itself.

Once the mobile node gets confirmation that the home agent has received the request, the mobile node shuts down, steps 15 through 19 in FIG. 1.

When the mobile node rejoins the network, it asks the home agent to stop responding on its behalf, (a “recovery request”) step 30, and resumes communication by opening a window. The agent then stops handling the remote peer’s persist probes, and removes the

binding that had been requested, step 29. The remote peer exits persist state when it receives a window advertisement greater than 0 (opening of the window).

Alternatively, the home agent may be configured with a timeout, such that if it does not receive a recovery request from the mobile node within the timeout, it can terminate the 5 request for handling the remote peer's probes. It is preferable to contact the nearest agent to the mobile node, since it is likely to be on the same link or a short hop away. Thus, the invention can be further modified to pass the persist probe handling to the foreign agent.

## Extension of TCP

10 As an alternative to the preferred embodiment of the invention, a client can request the remote peer to support a persistent TCP connection, based on two additional TCP options. This alternative requires neither an intermediate agent, nor repeated persist probes from the remote server. However, it does require modification of the TCP stack to support 15 the options.

### 1) TCPOPT\_PERSIST\_REQ Option:

Referring to FIG. 3, the remote peer is requested to support the persist timeout that will be requested later. If the remote peer supports the option, the client will use the 20 TCPOPT\_PERSIST\_TO option to inform it of the time for which the connection may be inactive. The remote end will keep the connection open, but will not probe it until the time period requested times out. TCPOPT\_PERSIST\_REQ option is sent with the SYN segment, step 32. It includes a cookie value that must change with every SYN segment, including retransmit.

If the remote endpoint is willing to accept TCPOPT\_PERSIST\_TO option later, the 25 remote endpoint ACKs the SYN with TCPOPT\_PERSIST\_REQ, step 33, but without the cookie. The mobile node records the fact that the remote endpoint has accepted the persist request option, and the connection is setup normally, step 34. If the ACK does not have the option included, then the request has been denied.

### 2) TCPOPT\_PERSIST\_TO Option:

When sending the window advertisement of 0 at the time of shutdown, the client may 30 add this TCP option, step 36. It includes the cookie that had originally been sent with the

SYN segment for verification by the remote endpoint. This option specifies the timeout period in which the mobile host is expected to reconnect.

The remote node is required to acknowledge the receipt of this window advertisement with an immediate persist probe, step 38. This is regarded as an acknowledgment by the 5 mobile node. Notably, a mobile node might retransmit the window advertisement multiple times (a retransmit is possible if the node does not get the persist probe), in which case the server is required to reply each time. When the mobile node receives the acknowledgment, it shuts down, step 39.

The server will not probe further until the timeout specified in the 10 TCPOPT\_PERSIST\_TO option. If the server does not receive any response when the timeout expires, it will retransmit in the usual way, *i.e.*, it will start the persist probes, just as it would have if it had received a window advertisement of zero bytes and there was no TCPOPT\_PERSIST\_TO option. The server will eventually terminate the connection if it does not receive any reply from the mobile node.

15 If the mobile node rejoins the network before the expiry of TCPOPT\_PERSIST\_TO timeout or before the remote server gives up persist probes, it will send a window open advertisement to the remote server, and the communication will continue normally.

The above two alternatives may be used in combination with the preferred 20 embodiment, or in the alternative, one or both of the options can be selectively used by the administrator without any impact on the working of any other protocol or network.

### Alternative Embodiments

In an additional alternative embodiment of the present invention, the invention could 25 be implemented for non-mobile clients. Essentially, this is accomplished by handing off the handling of the remote endpoints persist probe to a device that is not going down.

In yet another alternative embodiment of the present invention, periodic saves of the 30 TCP endpoint state could be made to protect against a possible crash of the system. The application must be capable of crash recovery. Then, the connection could be resumed on reboot.

In still another alternative embodiment, the invention can be used to migrate connections between server systems. This method is useful when a system has to be shutdown for maintenance or upgrades. The state can be saved on a medium accessible to the system that is going to take over, or it could be sent to the system over the network. The 5 system being shutdown advertises a window of zero, causing the remote system to go into persist state. The system taking over the connection replaces the home agent in this solution, and it will, as required, also take over the IP address of the system being shutdown. The original system then shuts down, at which point the system taking over the connection restarts the application and reacquires the connection as described in the invention. The new 10 system must adapt slightly to the prior handling of the persist requests. If TCP extensions are used, then the cookie needs to be passed to the server taking over service. However, in this case the server taking over the endpoint will not have to field persist probes.

Moreover, although the present invention is a method to keep TCP connections alive across reboots, the invention is not limited to TCP. An alternative embodiment for UDP 15 could be conducted analogously as far as the application is concerned. UDP deactivation is much simpler, and there is no protocol state to be maintained.

For the orderly shutdown of the UDP endpoint (Deactivation), the UDP engine fails/blocks all calls from the application. All data in the socket buffers would be dropped, and the local endpoint state is transferred to the fail-over standby. The state information thus 20 comprises the local IP address, local port, the foreign address and the foreign port. The socket and the UDP control block would be dissociated from the file descriptor. Then, the rest of the application recovery and start would be as described in the case of TCP applications.

Computer-readable signal bearing media include, but are not limited to, floppy disks, 25 hard disks, tape and CD-ROMs.

It will be appreciated that, although preferred and alternative embodiments of the invention have been described herein for purposes of illustration, various modifications may be made without departing from the spirit and scope of the invention. Therefore, it is manifestly intended that this invention be limited only by the claims and the equivalents 30 thereof.